

APPLICATION FOR U.S. LETTERS PATENT

TITLE:

TRANSACTION SCHEDULING FOR A BUS SYSTEM
IN A MULTIPLE SPEED ENVIRONMENT

INVENTORS:

John I. Garney
John S. Howard

PREPARED BY:

KENYON & KENYON
333 W. SAN CARLOS STREET
SUITE 600
SAN JOSE, CA 95110

(408) 975-7500

TRANSACTION SCHEDULING FOR A BUS SYSTEM
IN A MULTIPLE SPEED ENVIRONMENT

Field of the Invention

The present invention pertains to the field of data communication in a digital system.

5 More particularly, the present invention relates to host controllers and hubs used to transfer information on a bus.

Background Of The Invention

A computer or similar device typically has a bus that connects peripherals to the computing system. Sometimes a hub or multiple hubs may be placed in between the peripheral and the computing system (host). A hub provides a fan-out capability by allowing multiple peripherals to be connected to the hub which is in turn connected to the host or a daisy-chain of hubs one of which is ultimately connected to the host. Some of the peripherals operate at a high data rate and some operate at a low data rate. Due to a variety of advances (e.g., computing power) in computers (hosts) and peripherals, the data rates at which some peripherals operate has increased significantly. The increase in data rates cannot be met using existing bus standards. For example, the relative difference between the highest data rate peripheral on a bus and the lowest data rate peripheral on a bus has increased to the point that existing solutions for allowing high data rate peripherals and low data rate peripherals to co-exist on the same bus are typically not very efficient. Additionally, existing solutions for

allowing hosts to communicate with both advanced, high data rate devices and legacy, low data rate devices usually require the host and/or hub to be relatively complex and costly.

The increased demand for high data rates, as described above, cannot be met using existing buses nor using the bus architecture and protocols of existing buses. For example,
5 the Universal Serial Bus (USB) Specification Revision 1.1, September 23, 1998, (USB Standard) is limited to a full-speed data rate of 12 Mb/s (megabits per second) and a low-speed data rate of 1.5 Mb/s. Examples of relatively high data rate peripherals, include cameras, compact disc players, speakers, video cameras, microphones, video display devices, and scanners among other devices. Unfortunately, many of these devices have data rate
10 requirements that exceed the data rates supported by USB. For example, a video display device can have a data rate in excess of 20 Mb/s.

Existing solutions for allowing high data rate peripherals and low data rate peripherals to co-exist on the same bus are typically not very efficient when used for buses whose ratio of the highest data rate supported on the bus to the lowest data rate supported on the bus is relatively large. Examples of low data rate peripherals, include mice and joy-sticks that need to co-exist along with the high data-rate peripherals. A mouse typically has a data rate significantly less than 0.1 Mb/s. When the ratio of the highest data rate to the lowest data rate is relatively small, solutions such as speed-shifting and non-multiplexed store-and-forward transactions are tolerable despite their relative inefficiency.
15

In USB, for example, speed-shifting refers to a host communicating at a low-speed with low data rate peripherals and alternatively at full-speed with high data rate devices (speed-shifting). Unfortunately, the amount of data actually transmitted over the bus
20

(effective throughput) is less than that achievable by limiting the bus to full-speed transactions. Speed-shifting is also employed by "Firewire" or Institute of Electrical and Electronics Engineers (IEEE) Standard 1394 for a High Performance Serial Bus, 1995. Even though IEEE 1394 supports multiple data rates, up to 400 Mb/s, speed-shifting and the relatively high cost of Firewire make it an undesirable technology. The inefficiency of Firewire can be relatively severe when speed-shifting occurs in communicating between a 0.1 Mb/s mouse and a 20 Mb/s video display device.

Non-multiplexed store and forward transactions occur when a host (1) transmits at a high data rate a packet to a store-and-forward hub, (2) waits for the hub to forward at the low data rate the packet to the peripheral, (3) waits for the peripheral to respond at the low data rate to the hub, and (4) receives from the hub at a high data rate the peripheral's response to the packet. When the ratio of the highest data rate supported on the bus to the lowest data rate supported on the bus is relatively large, this co-existence solution may also result in a low effective throughput or bandwidth because of the time wasted in waiting for the hub to forward the packet at the low data rate and for the peripheral to respond at the low data rate.

As described above, existing buses are not capable of providing the high data rates required by modern peripherals. Additionally, existing solutions that allow high data rate peripherals and low data rate peripherals to co-exist on the same bus typically result in the bus having inefficient performance. Consequently, it is desirable to provide the high data rates required by modern peripherals, and efficient solutions allowing high data rate devices and low data rate devices to co-exist on the same bus.

Summary of the Invention

According to an embodiment of the invention a method for communicating data using a hub is described. The method includes determining a first estimated unused capacity left in a first frame in which a second transaction is to be performed between a hub and an agent.

5 The method then includes determining an amount of a first data that can fit into the estimated unused capacity and that is to be sent to the hub during a first transaction and then sent by the hub to the agent during the second transaction. The method also includes sending the first data to the hub during the first transaction.

Brief Description of the Drawings

10 The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which like references denote similar elements, and in which:

Figure 1a illustrates a block diagram of a digital system using a protocol in accordance with the present invention;

15 **Figures 1b, 1c & 1d** each illustrates a process showing a method in accordance with this invention for communicating between a host controller and a hub;

Figure 1e illustrates a hub in accordance with the present invention;

Figures 2a & 2b illustrate state machine diagrams for a host controller and a hub, respectively, performing a transfer in accordance with this invention;

Figures 3a & 3b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention;

Figures 4a & 4b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention;

5 Figures 5a & 5b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention;

Figures 6a & 6b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention;

10 Figures 7a & 7b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention;

Figures 8a & 8b illustrate best case and worst case frames for data transfers in accordance with the present invention;

Figure 9 illustrate timing diagrams for host controller-hub transactions and hub-agent transactions; and

15 Figure 10 illustrates a host controller's budgeted hub-agent frame, a host controller-hub frame, and an actual hub-agent frame.

Detailed Description

A method and apparatus for communicating between a host and a peripheral (agent) is described, where the agent communicates at a different speed and/or protocol than the host. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced in a variety of bus systems, especially serial buses, without these specific details. In other instances well known operations, steps, functions and devices are not shown in order to avoid obscuring the invention.

Parts of the description will be presented using terminology commonly employed by those skilled in the art to convey the substance of their work to others skilled in the art, such as device or controller drivers, bus or host controllers, hubs, bus agents or agents, and so forth. Also, parts of the description will also be presented in terms of operations performed through the execution of programming instructions or initiating the functionality of some electrical component(s) or circuitry, using terms such as, performing, sending, processing, packaging, scheduling, transmitting, configuring, and so on. As well understood by those skilled in the art, these operations take the form of electrical or magnetic or optical signals capable of being stored, transferred, combined, and otherwise manipulated through electrical components.

Various operations will be described as multiple discrete steps performed in turn in a manner that is most helpful in understanding the present invention. However, the order of description should not be construed as to imply that these operations are necessarily performed in the order that they are presented, or even order dependent. Lastly, repeated

usage of the phrases "in one embodiment," "in an embodiment," "an alternative embodiment," or "an alternate embodiment" does not necessarily refer to the same embodiment, although it may.

Figure 1a illustrates a block diagram of a bus using a protocol in accordance with the present invention. Bus 100 includes a system 102 having a host controller 110 which is coupled to hub 120 which is in turn coupled to agent 130. Host controller 110 has an associated device driver 108 that executes on system 102. Examples of agents include cameras, compact disc players, speakers, microphones, video display devices, scanners, and joy-sticks and mice, among other devices. System 102 can include any digital system capable of digital communication, especially laptop computers, desktop computers, servers, set-top boxes, entertainment systems, and game machines. Consequently, this invention can be practiced with a variety of digital devices using digital communication.

Two arrows 101a and 101b are drawn in Figure 1a to provide a frame of reference as to the direction of communication among the host, hub and agent. The direction from the host to the hub and on to the agent is indicated by arrow 101a and is referred to as the downstream direction or downstream (out). The direction from the agent to the hub and on to the host is indicated by arrow 101b and is referred to as the upstream direction or upstream (in).

Host controller driver 108 facilitates communications or transactions along bus 100 (e.g., on behalf of an application executing on system 102) by processing packets of information destined for agent 130 and scheduling the packets for transmission by controller 110. Host controller 110 sends data to and receives data from agent 130 via hub 120. Agent 130 communicates at a different or lower data rate (agent data rate) than the data rate of

host controller 110 (host data rate). While only one agent is shown coupled to hub 120, it should be apparent that additional agents (not shown) can be attached to hub 120 or to other hubs (not shown). These additional agents may communicate at the host data rate or the agent data rate. Furthermore, while agent 130 is shown coupled to hub 120, it may be coupled to
5 hub 120 through at least one conventional repeater type hub that operates at the agent data rate. A conventional repeater type hub repeats signals it receives on its upstream side on its downstream ports, and vice versa. The conventional repeater type hub may in turn have one or more agents 130 attached to it.

Host controller 110 and agent 130 have a master-slave relationship which means that
10 the host initiates typically all data transactions between the host and an agent; i.e., the agent only responds to requests from the host but never initiates a transaction. Hub 120 has store-and-forward buffers (not shown) that allow hub 120 to temporarily store downstream information received from host controller 110 and destined for agent 130, and to temporarily store upstream information received from agent 130 and destined for host controller 110.

15 Since agent 130 and host controller 110 communicate at different data rates, it is desirable to enhance the effective throughput on the bus by providing a protocol that would allow host controller 110 to both (1) communicate at its higher data rate and (2) not have to wait for responses from agent 130 before engaging in another transaction. The protocol of the present invention allows host controller 110 to take advantage of the store-and-forward
20 characteristic of hub 120 to allow the host controller 110 to communicate at its higher data rate and to engage in another transaction instead of waiting for a response from agent 130, if a response is required. The protocol of the present invention also provides robustness and

reliability to transactions performed between controller 110 and hub 120. Additionally, the host controller and/or hub of the present invention allow increased effective throughput on the bus and provide increased responsiveness to agents (not shown) that communicate at the host data rate and that are attached to hub 120 or other hubs.

5 Figure 1b illustrates a process 150 showing a method in accordance with this invention for communicating with an agent having a lower (or different) data rate, than the data rate of a host controller. The agent may also have a different protocol than the host controller. Process 150 can be used to effect a variety of information transfers between host controller 110 and agent 130. For ease of understanding process 150 will only be described here with regards to a bulk out transfer. However, process 150 can be used with other information transfers described herein, below. In a bulk out transfer data is transferred from host controller 110 to agent 130 via hub 120. The bulk out transfer is defined according to an embodiment of this invention as an asynchronous transfer type. However, it should not be concluded from this definition that any bulk and/or out transfer need be asynchronous.

10

15 At step 152 in process 150 a start split transaction is performed. The start split transaction communicates downstream information from host controller 110 to hub 120. Some of the downstream information communicated to hub 120 is temporarily buffered in hub 120. The buffers in hub 120 largely behave in a first-in-first-out (FIFO) manner, and are described in greater detail below. Some time after the downstream information is buffered, 20 hub 120 performs a hub-agent transaction (not shown) with agent 130 based on some of the buffered downstream information. The relative timing of the hub-agent transaction need not be described herein because one of ordinary skill in the art would recognize that this is an

application or implementation detail for which there are many possibilities. The hub-agent transaction may result in upstream information being buffered in hub 120. Some time after the downstream information is buffered, at step 156 a complete split transaction is performed. The complete split transaction communicates buffered upstream information from hub 120 to host controller 110. The relative timing of the complete split transaction need not be described herein because one of ordinary skill in the art would recognize that this is an application or implementation detail for which there are many possibilities.

A benefit of the split transaction protocol is that it allows controller 110 to initiate communication (start-split transaction) with agent 130 , engage in another function, or engage 10 in another communication with another agent (low data rate or high data rate agent) at step 154, and then return to complete the communication that was initiated earlier with the low data rate agent. By communicating using split-transactions, controller 110 communicates at high data rates without speed-shifting and does not sit idle while waiting for hub 120 to communicate with agent 130. The time that would have been spent being idle can be used to 15 communicate with another agent. In an alternative embodiment in accordance with the present invention, controller 110 may engage in speed-shifting with some agents while engaging in split-transaction communication with other agents.

The ~~start~~ split and the complete split transactions (split transactions) described above may be used to ~~implement~~ a variety of transfer types (e.g., read or write) for communicating data between ~~controller~~ 110 and agent 130. In an embodiment of this invention four transfer types (or transfer requests) are defined: bulk out/in, control out/in, interrupt, isochronous. It should be apparent to one of ordinary skill in the art that the spirit and scope of this invention

includes other embodiments with fewer, more or different transfer types. Each of the transfer types provides different levels of robustness, reliability, synchronization, asynchronous operation, error detection and correction of the communication flow, and other characteristics that should be apparent to one of ordinary skill in the art. For example, bulk out/in provides

5 large asynchronous data transfers from controller 110 to agent 130 or in the opposite direction. Control out/in also provides asynchronous data transfer from controller 110 to agent 130 or in the opposite direction, but the data is typically control information used to control the operation of elements (e.g., a tape drive) in agent 130 or system 100. Interrupt provides a periodic data transfer from controller 110 to agent 130 or in the opposite direction. If the

10 transfer is not successful, controller 110 may try again in an embodiment in accordance with this invention. Isochronous transfer provides a data transfer once every predetermined time interval. According to an embodiment of the present invention, the transfer may occur at any time during the time interval. If the transfer is not successful, controller 110 will not repeat the transfer. In an alternative embodiment in accordance with the present invention, the

15 isochronous transfer may provide for repeat transfers.

The split transactions may include a number of phases depending on the transfer type being implemented. Each of the split transactions may have up to three phases: token, data, and handshake. However, depending on the transfer being performed, some transactions may have fewer phases. In an embodiment of the present invention, bulk and control can use the same phases in each of their respective split transactions. The phases for each of the transfer types described above are shown in Table 1, below. Presence of an "X" in a cell of the table indicates that the split transaction for the transfer type has the phase indicated at the top of the

column in which the cell resides. While in this embodiment the token and data phases are separate for each of the transfer types, in alternative embodiments the token and data phases may be combined. It should be apparent that in alternative embodiments transfer types may have fewer, more, or even different phases than those shown in Table 1 without departing 5 from the scope and spirit of the present invention.

Table 1
Start-Split Transaction Complete-Split Transaction

Transfer Type	Token	Data	Handshake	Token	Data	Handshake
Bulk-Control Out	X	X	X	X		X
Bulk-Control In	X		X	X	X	X
Interrupt Out	X	X		X		X
Interrupt In	X			X	X	
Isochronous Out	X	X				
Isochronous In	X			X	X	

Figure 1c illustrates in greater detail a process 160 showing a start split transaction for 10 a bulk out transfer in accordance with an embodiment of this invention. At step 162 a token packet including hub identification information, agent and endpoint identification 15 information, transfer type, indicator for specifying direction of transfer (in or out), and data rate identification is sent from host controller 110 to hub 120. Hub identification information, and agent and endpoint identification information, and direction are together commonly 20 referred to here as transaction addressing information. The agent identification information identifies the particular agent with which the host is attempting to communicate. The endpoint identification information identifies a particular portion in the agent with which the

host is attempting to communicate. Examples of endpoints include: left speaker and right speaker of a speaker hub, or speaker and microphone of telephone handset. The transfer type in the transaction addressing information is not limited to the types described herein (e.g., bulk out, interrupt, isochronous, control), but can include other types known in the art without departing from the spirit and scope of this invention. Data rate identification specifies the data rate with which the hub-agent transaction described in connection with process 150 above will be performed. For an embodiment in which the hub-agent transaction is performed in accordance with the USB standard, data rate identification will specify either 12 Mb/s (full-speed) or 1.5 Mb/s (low-speed). At step 164, a data packet is sent from host controller 110 to hub 120. At step 166, a first acknowledgement is received by host controller 110 from hub 120, if the data packet was decoded properly by hub 120. The first acknowledgement indicates whether the data was decoded properly by hub 120 or whether hub 120 wants to be tried later (e.g., hub 120 had full buffers and was not able to accept the data).

Figure 1d illustrates in greater detail a process 170 showing a complete split transaction for a bulk out transfer in accordance with an embodiment of this invention. At step 172, a second token packet including transaction addressing information is sent from the host to the hub. At step 174, a second acknowledgement is received by host controller 110 from hub 120, where the second acknowledgement can either (1) include handshake information received by hub 120 from agent 130 during the hub-agent transaction described above in connection with Figure 1b or (2) indicate that hub 120 does not yet have information based on the hub-agent transaction to forward to host controller 110 (e.g., the hub-agent transaction has not yet been completed). The handshake information indicates whether (1)

agent 130 properly received data during the hub-agent transaction (ACK), (2) agent 130 indicated that it is not able to operate normally (STALL), or (3) agent 130 indicated that it wanted to be tried later (NAK). While the first and second acknowledgements and the handshake information have been described as specifying certain indicators, it should be apparent to one of ordinary skill in the art that these acknowledgements and handshakes and other ones described herein may represent other indications. Additionally, acknowledgements and handshakes different from or additional to the ones described herein may be added in an alternative embodiment without departing from the spirit and scope of the invention.

While the above description has generally been presented in the context of agent 130 and hub 120 communicating at a lower data rate than the data rate between hub 120 and host controller 110, those skilled in the art will appreciate that the present invention may be practiced to bridge a lower data rate to a higher data rate instead, or even equal data rates but different protocols.

While in **Figure 1** only one hub was shown in between the agent and the host there can be multiple hubs between any particular agent and the host. While only six transfer types have been described, those skilled in the art will appreciate that other types can be used without departing from the scope or spirit of this invention.

Each of figures 2a, 2b, 3a, 3b, 4a, 4b, 5a, 5b, 6a, 6b, 7a & 7b illustrates a state machine diagram for performing a transfer using a host controller and a hub in accordance with this invention. Figures with an "a" as a suffix show the state machine diagram for a host controller; the state machine may be performed on host controller 110 described above in connection with **Figure 1a**. Figures with a "b" as a suffix show the state machine diagram for a hub; the state machine may be performed on hub 120 described above in connection with **Figure 1a**. The state machines illustrated in these figures show processes having multiple steps. It should be apparent that some of the steps may be divided into more numerous steps or combined into fewer steps without departing from the scope and spirit of this invention.

The state machines are not described in great detail because their operation should be apparent to one of ordinary skill in the art. For ease of understanding, the processes performed by the state machines are not described separately. Since the processes performed by the state machines operate in tandem (i.e., each process has steps whose execution depends on the occurrence of events or steps in the other process), the description of the processes are interleaved to facilitate understanding.

Figures 2a & 2b illustrate state machine diagrams for a host controller and a hub, respectively, performing a transfer in accordance with this invention, specifically a split out bulk/control transfer. Process 200 and process 260 show the state machine for a host controller and a hub, respectively. Process 200 includes start split transaction, having a token phase (XOUT) and a data phase (DATAx), which may be repeated up to three times by the host controller when timeouts occur during the phases of a transaction attempt between the

host controller and the hub. In response to the start split transaction, process 260 will either propagate through states that will accept the data (ACK), respond to a host controller retry after a communication failure of a hub handshake to the host controller (ACK, don't accept data), request a host controller retry due to lack of space to hold the start transaction

5 information (NAK), or timeout up to three times. Process 200 shows the host controller response to a complete split transaction (XIN) when the transaction between the hub and agent was successfully processed (Advance), resulted in a NAK from the agent (NAK), received an indication that the agent was not able to function normally (STALL), was not yet completed by the hub or agent (NYET), or the XIN or its response had some communication failure and

10 resulted in up to three timeouts between the host controller and hub. In response to the complete split transaction, process 260 will either indicate that the transaction between the hub and the agent has not finished (NYET) or will provide an indication (ACK, NAK, STALL) of what transpired during the transaction between the hub and the agent.

Figures 3a & 3b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention, specifically a split out interrupt transfer. Process 300 and process 360 shows the state machine for a host controller and a hub, respectively. Process 300 includes a start split transaction, having a token phase (X^{OUT}) and a data phase (DATAx), which is not repeated by the host controller because according to an embodiment the interrupt out transfer is time sensitive and need not be repeated if it is not successful on the first try. In response to the start split transaction, process 360 will either accept the data (ACK), or do nothing. Process 300 shows the host response to a complete split transaction (XIN) when the transaction between the hub and agent

15

20

was successfully processed (Advance), received an indication that the agent was not able to function normally (STALL), was not yet completed by the hub or agent (NYET), or the XIN or its response had some communication failure and resulted in up to three timeouts between the host controller and hub. In response to the complete split transaction, process 360 will either indicate that the transaction between the hub and the agent has not finished (NYET) or will provide an indication (ACK, NAK, STALL, NYET) of what transpired during the transaction between the hub and the agent.

Figures 4a & 4b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention, specifically a split out isochronous transfer. Process 400 and process 460 shows the state machine for a host controller and a hub, respectively. Process 400 includes a start split transaction, having a token phase (XOUT) and a data phase, neither of which is repeated. Process 400 does not include a complete split transaction according to an embodiment. Process 460 allows the hub to agent transaction data to be subdivided into multiple sections to minimize buffering required in the hub; the host controller can send the next section of data just before the hub needs to send it to the agent. In this case each split start transaction is marked with ALL, BEGIN, MIDDLE or END so that the hub can detect when a communication failure has caused it to not receive a data section in the correct sequence. In response to the start split transaction, process 460 will accumulate a data payload having one data section (ALL), two data sections (BEGIN, END), or three or more data sections (BEGIN, MIDDLE...MIDDLE, END).

Figures 5a & 5b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention, specifically a split in bulk/control transfer. Process 500 and process 560 shows the state machine for a host controller and a hub, respectively. Process 500 includes start split transaction, having a token phase (XOUT), which is repeated up to three times by the host controller when timeouts occur during transaction attempts between the host controller and the hub. In response to the start split transaction, process 560 will either acknowledge and accept the transaction (ACK, accept xact), respond to a host controller retry after a communication failure of a hub handshake to the host controller (ACK, ignore xact), or request a host controller retry due to a lack of space to hold the start transaction information. Process 500 shows the host controller response to a complete split transaction (XIN) when the transaction between the hub and agent was successfully processed (Advance), received an indication that the endpoint is unable to function normally (STALL), was not yet completed by the hub or agent (NYET), ignores data received from the hub because it is corrupted, or the XIN or its response had some communication failure and resulted in up to three timeouts between the host controller and hub. In response to the complete split transaction, process 560 will either indicate that the transaction between the hub and the agent has not finished (NYET) or will provide an indication (NAK, STALL) of what transpired during the transaction between the hub and the endpoint, or send the data received from the agent by the hub to the host controller.

Figures 6a & 6b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention, specifically a split in interrupt transfer. Process 600 and process 660 shows the state machine for a host

controller and a hub, respectively. Process 600 includes a start split transaction, having a token phase (XOUT). In response to the start split transaction, process 660 will accept the transaction (accept xact). Process 600 shows the host response to a complete split transaction (XIN) when the transaction between the hub and agent was successfully processed (Advance),

5 received an indication that the endpoint is not able to function normally (STALL), was not yet completed by the hub or agent (NYET), received a NAK from the agent (NAK), retries token phase if data received is an agent retry of the previous transaction request (ignore data), or the XIN or its response had some communication failure and resulted in up to three timeouts before the host controller gives up communicating with the agent. In response to the complete

10 split transaction, process 660 will indicate that a timeout occurred when the hub was communicating with the agent (NYET), or the hub didn't receive the start transaction for this request and had no corresponding response information (STALL) or provide an indication (NAK, STALL) of what transpired during the transaction between the hub and the agent, or send the data to the host controller.

15 Figures 7a & 7b illustrate state machine diagrams for a host controller and a hub, respectively, performing another transfer in accordance with this invention, specifically a split in isochronous transfer. Process 700 and process 760 shows the state machine for a host controller and a hub, respectively. Process 700 includes a start split transaction, having a token phase (XOUT). In response to the start split transaction, process 760 will accept the transaction (accept xact). Process 700 shows the host response to a complete split transaction (XIN) when the transaction between the hub and agent was successfully processed and all the data has been returned (TAadvance) or there is more data to return (DAadvance) or an error

20

occurred (record error) due to a communication failure between the host and hub (timeout or bad cyclic redundancy check) or the agent (NAK) or a hub problem with the agent (NYET).

In response to the complete split transaction (XIN), process 760 will indicate that data received from the agent had a bad cyclic redundancy check (NAK) or the agent didn't respond

- 5 (NYET) or the hub had no information about this complete-split, or send the data to the host controller either indicating all the data has been returned (DATA0) or more data is to be returned (MDATA).

It is useful at this point to summarize the description of the above protocol before describing the remaining apparatus and methods of the present invention. The protocol

- 10 described above allows a host controller to transfer data to or receive data from an agent via a hub. The protocol allows the host controller to engage in a first transaction (start split transaction) in which a transfer request is communicated to the hub. After the host controller performs the first transaction, it may engage in an intermediate transaction with the same hub or another hub without waiting for the hub and agent to perform the transfer request (i.e., engage in the transfer of data to or from the agent). The intermediate transaction may include a transfer request for the agent, another agent on the same hub as the agent, or another agent on yet another hub. After the hub engages in the transfer of data to or from the agent, the host controller performs a complete split transaction (or second transaction) to get the result (e.g., data or handshake sent from the agent to the hub) of the transfer performed between the hub
15 and the agent. By allowing the host controller the capacity to engage in an intermediate transaction, instead of waiting for the hub to perform the transfer request (or third transaction) with the agent, the effective throughput of a bus using a protocol in accordance with this
20

invention can be significantly greater than buses which involve speed-shifting or which require the host controller to wait for the hub to perform the transfer with the agent before initiating another transfer.

While the above protocol defines the sequence of transactions involved in

5 communicating data across a bus, it does not explicitly describe the timing for the transactions or transfers which will result in data being sent to or received from an agent. However, timing of transfers for agents is important because agents typically require data to be sent to or received from the host controller on a periodic (e.g., isochronous or interrupt) or asynchronous (e.g., bulk or control) basis. Additionally, while the above protocol allows a host controller to

10 perform an intermediate transaction (or even multiple intermediate transactions) between the start split and the complete split transaction, the above protocol does not explicitly describe how transfer requests are stored in the hub and how the hub and agent perform transfer requests without requiring the host controller to wait for the transfer request to be performed before engaging in an intermediate transaction. The issues not addressed by the above

15 description of the protocol, namely the timing of transfer requests and the processing (i.e., buffering and performance) of transfer requests by a hub, are addressed by the following descriptions of methods and apparatus in accordance with the present invention. The present invention includes a method and apparatus for scheduling transfers of data to and from an agent and a method and apparatus for processing transfer requests at a hub. The method and apparatus for scheduling transfers of data is described first and then the method and apparatus for buffering and performing transfer requests is described second.

20

Referring again to Figure 1, the process of scheduling transfers of data first begins

when system 102 performs configuration for agents attached to the bus of the system. Configuration can occur upon system initialization or boot-up or when an agent is attached to the bus after initialization. The scheduling of data transfers depends on the transfer type associated with an agent (or an endpoint of an agent) and the amount of data involved in the transfer. The manner in which agents having associated periodic transfers, such as isochronous and interrupt, are handled is described first, below, and then the manner in which agents having associated asynchronous transfers, such as bulk and control, are handled is described next.

During the process of configuration, each endpoint of each agent informs the system of the endpoint's associated maximum data payload size and transfer type (e.g., isochronous in/out, interrupt, bulk, control). The manner in and device by which each agent informs the system is well understood by those of ordinary skill in the art and need not be described here. The maximum data payload size is the largest amount of data that a transfer to or from an agent will entail. The system relays the data payload size and the transfer type to the host controller. The manner in which the system relays the size and transfer type to the host controller is well understood by those of ordinary skill in the art and need not be described here. The host controller uses the aforementioned two pieces of information associated with each endpoint to generate a budget list for the periodic transfers of the endpoints. In alternate embodiment, a software driver such as the host controller driver or even a hardware driver can generate the budget list and perform the scheduling operations described below. The budget list gives the earliest time that a transfer (sending or receiving a data packet) may occur as well as the latest time that a result associated with the transfer may be available. The earliest

time that a transfer can occur depends on the amount of time taken up by each of the previous transfers assuming that the previous transfers happened under the best of circumstances (defined below). The latest time that a result associated with a transfer may be available depends upon the amount of time taken up by each of the previous transfers assuming that 5 previous transfers happened under the worst circumstances (defined below) and the time required for the transfer under the worst circumstances. The earliest time that a transfer can occur is important because the host controller needs to have the transfer request buffered at the hub before that time so that as soon as the hub finishes the transfer request that is ahead of the buffered request the hub can turn its attention to the buffered transfer request. The latest 10 time that a result associated with the transfer may be available is important because after the latest time it is substantially certain that the result will be available for the host controller to retrieve. If the host controller were to try to retrieve the result from the hub before the latest time, meaning that the result is not available yet, a less efficient protocol employing multiple retrievals would be required.

15 Transfers happen under the best of circumstances when each transfer involves the maximum data payload and there is substantially no bit-stuffing. Transfers happen under the worst of circumstances when each transfer involves maximum data payload and there is maximum bit-stuffing. Bit-stuffing occurs according to an embodiment of the present invention because the signals on the bus obey non-return-to-zero (NRZ) signaling. According 20 to an embodiment of the present invention, bit-stuffing may increase the size of the maximum data payload by 16%. While in an embodiment, the best circumstances and the worst circumstances have been defined in terms of bit-stuffing, it should be apparent that in

alternative embodiments the circumstances can be described in terms of other things that can expand or decrease the size (or time) of transfers, or generate delays.

The generation of a budget list in accordance with the present invention will now be described. While one way of generating a budget list is described herein, it should be apparent that the spirit and scope of the present invention encompasses other possible ways.

A budget list is a list of the allowable periodic transactions that can occur in a particular frame template. A frame is a continuously occurring period of time defined as part of the bus specification that is sufficient to provide for one or more transactions. In an embodiment, a 1 millisecond time period is defined for a frame. A different budget list is constructed for each frame that has a different set of allowable periodic transactions. A frame template is a description of a particular periodically repeating frame that provides for some maximum number of transactions, each transaction having some maximum data payload. A frame contains some number of transactions of some actual data payload, while a frame template describes a potential budgeted frame. Each budget list has an associated best case information and an associated worst case information. The best case information describes the situation in which each transaction in the frame template occurs under the best circumstances. In an alternative embodiment, transfers rather than transactions are represented by the best case information and worst case information.

Figure 8a illustrates a diagram for a best case frame template 800 in accordance with the present invention. Block 805 represents the transfer associated with a first endpoint, where the transfer happens under the best case. Block 810 represents the transfer associated with a second endpoint, where the transfer happens under the best case. The remaining blocks

815-835 represent similar best case transfers for other endpoints configured by the system.

The worst case information describes the situation in which each transfer in the frame template occurs under the worst circumstances. Figure 8b illustrates a diagram for a worst case frame template 850 in accordance with the present invention. Block 855 represents the transfer associated with the first endpoint, where the transfer happens under the worst case. Block 860 represents the transfer associated with the second endpoint, where the transfer happens under the worst case. The remaining blocks 865-885 represent similar worst case transfers for other endpoints configured by the system. It should be apparent that the relative sizes of the blocks are only meant for purposes of illustration.

10 The significance of best case frame template 800 and the worst case frame template 850 will now be explained by examining blocks 805, 810, 855 and 860. Blocks 805 and blocks 855 represent the best case and worst case transfer, respectively, for the first endpoint. The earliest that the transfer for the first endpoint can finish is Ta. The latest that the transfer for the first endpoint can finish is Tb. The foregoing suggests that the transfer associated with 15 the second endpoint can begin as early as time Ta or begin as late as time Tb. Blocks 810 and 860 represent the best case and worst case transfer times, respectively, for the second endpoint. The earliest that the transfer for the second endpoint can end is Tc. The latest that the transfer for the second endpoint can finish is Td. The foregoing suggests that the information necessary for the second transaction must be available to the hub before time Ta 20 and any result due to the first transfer is substantially certain to be available after time Tb. Furthermore, it should be apparent that scheduling of a transfer depends on the time required to perform each of the previous transfer requests. Additionally, it should be apparent that the

scheduling of the complete-split transaction will depend on the time to perform the present transfer request, if any, associated with a start-split transaction.

Frame template 800 and frame template 850 are frame templates that are representative of a frame between the hub and the agent (hub-agent or classic frame).

- 5 According to an embodiment of the present invention, the host controller and the hub have another frame that is a fraction of the size (microframe) of the classic frame. More specifically, in an embodiment, eight microframes are equivalent to a single classic frame. Since the host controller and the hub communicate using microframes and since a transfer request before reaching the hub starts as a start split transaction, it is useful to describe the 10 time for performing the start split transaction in terms of microframes. Similarly it is useful to describe the time for performing the complete split transaction in terms of microframes. With regards to the start time for a start split transaction, according to an embodiment, a start split transaction should occur one microframe before the microframe in which the transfer between the hub and agent can occur according to the best case frame. With regards to the start time 15 for a complete split transaction, according to an embodiment, a complete split transaction should occur one microframe after the microframe in which the transfer of the associated start split transaction can finish according to the worst case frame. According to an embodiment of the present invention, the host controller uses the best case frame and the worst case frame for each frame template to create a start-split vector and a complete-split vector for each endpoint in each frame template. The start-split vector contains a value indicative of the microframe in which the start split transaction for an endpoint should occur during a classic frame. According to an embodiment, start-split vector has values ranging from -1 to 6. The 20

complete-split vector contains a value indicative of the microframe in which the complete split transaction for an endpoint should occur during a classic frame. According to an embodiment, the complete-split vector has values ranging from 1 to 8.

When host controller 110 is tasked by host controller driver 108 to perform a transfer

5 in a particular frame for a particular endpoint, host controller 110 examines the best case vector that is associated with the endpoint and the frame template that corresponds to the particular frame. The value in the start-split vector for a given endpoint and frame template indicates the microframe at which a start split transaction is to be performed. Similarly, host controller 110 examines the complete-split case vector to determine the microframe at which a

10 complete split transaction is to be performed.

Figure 9 illustrates a sequence of hub-agent frames (or classic frames) 900 and a host controller-hub frame 950. Figure 9 illustrates the situation for the case where an isochronous out transfer requires no more than one microframe of activity between the host controller and the hub. The case where an isochronous out transfer spans two or more consecutive microframes is described below in connection with Figure 10. Frames 900 include frames 910-930. Frame 930 illustrates the host controller's budgeted frame for hub-agent activity for the isochronous out transfer. Frame 950 illustrates the timing of split transactions that occur between the host controller and the hub and that cause the hub-agent activity for the isochronous out transfer to be realized.

Frame 930 has a transfer 901 which is scheduled to start in microframe B and is scheduled to end in microframe B. In this illustration, transfer 901 is an isochronous out transfer, but it should be apparent that other transfers can be given a similar representation.

Frame 950 includes sub-frames 961 and 963. While transfer 901 is scheduled to occur 5 in hub 120 in microframe B and end in microframe B (i.e., in the hub-agent time frame represented by frames 900), the associated start-split transaction occurs in frame 961 and the associated complete-split transaction occurs in frame 963. In the host controller-hub time frame represented by frame 950, the associated start-split transaction occurs at time 951a sometime in sub-frame 961 and the associated complete-split transaction occurs at time 951b 10 sometime in sub-frame 963. Transfer 951 represents transfer 901 in the host controller-hub time frame. Other transfers and transactions can also occur in sub-frames 961 and 963. It should be apparent from **Figure 9**, that sub-frames 961-3 are available for other transfers and transactions with other agents.

While in the above description an isochronous out transaction has both a start-split 15 transaction and a complete-split transaction, it should be appreciated that in an alternative embodiment in accordance with the present invention an isochronous out transaction may not include a complete-split transaction. While the above description has generally been presented in the context of periodic transactions, the invention is not limited to periodic transaction but also ~~includes~~ asynchronous transfers such as the bulk and control transfers 20 described above. According to an embodiment of the present invention, ninety percent of a classic frame is set aside for periodic transfers. Asynchronous transfers will occur at the earliest time there is space available for them on the bus between the hub and the agent. The

hub will accept a transaction during a microframe when it has space available to hold the start-split transaction until it can be issued on the classic bus (i.e., between the hub and an agent).

Then the hub will issue the transaction during a classic frame when it has no other pending

periodic transactions. Once the hub has accepted a start-split, the host controller will later

5 issue a complete-split transaction to retrieve the results from the hub for that transaction on

the classic bus. The host controller driver issues an asynchronous split-transaction to the hub

whenever it has no other periodic transactions to issue during a microframe. The buffering

provided by the hub for bulk/control transactions is distinct and separate from the buffering

provided by the hub for periodic transactions for one embodiment.

10 **Figure 10** illustrates a host controller's budgeted hub-agent frame (or classic frame) 10-900, a host controller-hub frame 10-950, and an actual hub-agent frame 10-970. **Figure 10** illustrates activity (frame 10-950) between the host controller and the hub. **Figure 10** also illustrates activity (frame 10-970) between the hub and the agent for the case where an isochronous out transfer (10-972) spanning two or more microframes between the hub and the 15 agent is preceded by an isochronous in transfer (10-971) that finishes (time 10-971b) before the time it was supposed to finish (time 10-971c) according to the host controller's budgeted frame (10-900) for the isochronous transfer.

Frame 10-900 includes microframes 10-A, 10-B, 10-C. Frame 10-920 illustrates the host controller's budgeted frame for hub-agent activity for an isochronous in transfer 10-901 that precedes an isochronous out transfer 10-902. Frame 10-950 illustrates the timing of split 20 transactions that occur between the host controller and the hub and that cause the hub-agent

activity for the isochronous in transfer and the isochronous out transfer to be realized. Frame 10-970 illustrates the timing of the transfers that occur between the hub and the agent.

Frame 10-900 has a transfer 10-901 that is scheduled to start in microframe 10-B and is scheduled to end in microframe 10-B. In this illustration, transfer 10-901 is an isochronous in transfer, but it should be apparent that other transfer types are possible and can be given a similar representation. Frame 10-920 also has transfer 10-902 that is scheduled to start in microframe 10-B and is scheduled to end in microframe 10-C. In this illustration, transfer 10-902 is an isochronous out transfer, but other types of transfers are possible.

Frame 10-950 includes sub-frames 10-961, 10-962, 10-963, and 10-964. While transfer 10-901 is scheduled to occur in hub 120 in microframe B and end in microframe B (i.e., in the hub-agent time frame represented by frames 10-900), the associated start-split transaction 10-951 occurs in sub-frame 10-961 and the associated complete-split transaction 10-951* occurs in frame sub-10-963. Other transfers and transactions can also occur in sub-frames 10-961, 10-962, 10-963, and 10-964 with the agent. It should also be appreciated from Figure 10, that sub-frames 10-961, 10-962, 10-963 and 10-964 are available for other transfers and transactions with other agents.

While transfer 10-902 is scheduled to occur in hub 120 in microframes 10-B and 10-C, associated start-split transactions 10-952' and 10-952" each occur one microframe earlier than the microframe they are budgeted to occur in. Associated start-split transactions occur in sub-frames 10-961 and 10-962, respectively, and the associated complete split transaction 10-952* occurs in frame 10-964, a microframe later than the microframe in which transfer 10-902 is scheduled to be finished.

For one embodiment, the host controller determines the amount of bytes including both overhead bytes and data bytes that can be sent from the time that transfer 902 is scheduled to start (time 10-901b) in microframe 10-B until the end of microframe 10-B (estimated unused capacity). Overhead bytes include information sent in the token phase of a transaction. The amount of data that fits into the unused capacity after taking into account the capacity that is used up by the overhead bytes is referred to herein as the reduced data payload.

For one embodiment, host controller sends no more than the reduced data payload to the hub during associated start-split transaction 10-952'. By sending no more than the reduced data payload to the hub, the host controller conserves bandwidth on the downstream path between the host controller and the hub and the amount of buffering that is required in the hub.

Sending more than the reduced data payload takes up bandwidth on the downstream path that could have otherwise been spared for other high speed transactions with the hub or another hub or agent. Sending more than the than the reduced data payload may be wasteful whenever transfer 901 takes up its budgeted time.

While in the above description, the host controller sends no more than the reduced data payload, in an alternative embodiment the host controller sends additional data, but in no event does the host controller send a microframe's worth or more than a microframe's worth of data for the first start-split. For example, for one embodiment the host controller sends the reduced data payload and a fraction of the reduced data payload. The sum of the reduced data payload and the fraction of the reduced data payload is referred to herein as a fractionally overdriven payload. The fraction can be any percentage between 0% to the percentage that is representative of 100* (a microframe's worth of data less (one byte and any overhead

bytes))/(a microframe's worth of data). The largest percentage fraction herein refers to 100*
(a microframe's worth of data less (one byte and any overhead bytes))/(a microframe's worth
of data). For various embodiments of the present invention, the fraction may be 0%, 5%,
10%, 30%, 50%, 80%, 95%, the largest percentage fraction, and any number in between. It
5 should be appreciated that while a byte is used to count data other units such as word, bit,
nibble, or double word, and other units of measuring information or data are encompassed by
the present invention.

Since all the data needed for transfer 902 did not arrive in start-split transaction 10-
952, the hub should not start sending the reduced data payload received in start-split
10 transaction 10-952 to the agent unless the number of bytes that can be sent before the next
microframe starts is less than or equal to the sum of the reduced payload and the overhead
bytes. If the hub were to start sending data from the reduced data payload early, it is possible
for the hub to run out of data before the hub had received the additional data in start-split
transaction 10-952", causing an underrun condition in the hub-agent activity which is highly
15 undesirable. To prevent the hub from starting to send the reduced data payload to early, for
one embodiment, the host controller sends a multi-part indication in the token phase of first
associated start-split transaction 10-952'. The multi-part indication indicates to the hub that
the hub is not to start sending the reduced data payload to the agent in a microframe unless the
number of bytes that can be sent before the next microframe starts is less than or equal to the
20 reduced data payload plus any overhead bytes.

Frame 10-970 illustrates the situation where the hub does not take as much time to
perform transfer 901 as the host controller had budgeted for (i.e., transfer 10-971, which is

representative of transfer 901 in the hub-agent time frame, ends before its budgeted worst case end time) and the hub delays performing transfer 902 until the number of bytes that can be sent before the next microframe starts is less than or equal to the sum of the reduced data payload (or fractionally overdriven data payload) received during start-split transaction 10-
5 952' and any overhead data bytes needed to start the isochronous out transfer. Frame 10-970 includes microframes 10-981, 10-982, and 10-983 which illustrate activity between the hub and the agent associated with budgeted transfers 901 and 902. The hub, having received a start-split transaction at time 10-951a, performs an isochronous in transfer 10-971. The isochronous in transfer 10-971 is shown as being narrower than the budgeted isochronous in
10 transfer 901 because the isochronous transfer happened not to require all the time budgeted for the transfer (e.g., the peripheral or agent happened to have a smaller data payload to send to the host controller than the agent had asked for at the time of initialization). Even though the hub finished isochronous in transfer 10-971 early, the hub does not immediately go and start isochronous out transfer 10-972. Rather, for one embodiment, the hub examines the multi-part indication sent during start-split transaction 10-951'. If the hub determines that the transfer is one that spans several microframes, the hub waits to start isochronous out transfer 10-972 until the time left in microframe 10-982 is less than or equal to the time needed to send the sum of the reduced data payload (or fractionally overdriven data payload) received in
15 the start-split transaction that started at time 10-952a and any overhead bytes.

20 In an alternative embodiment, the hub determines whether a transfer is one that spans several microframes by examining a multi-part indication that is sent during initialization of the endpoints rather than during a start-split transaction. In this case, the multi-part indication

indicates that a particular out transfer associated with an agent likely or always spans two or more microframes. In this alternative embodiment, the hub will always delay the transfers associated with a particular endpoint or agent.

While in the above description an isochronous out transfer is the cause of the delayed performance of a transfer in the hub-agent time frame, it should be apparent that in alternative embodiments the delayed performance of a transfer may be due to other transfer types that are not isochronous or in transfers.

A method and apparatus for buffering periodic transfer requests and processing them in accordance with this invention will now be described. **Figure 1e** shows in greater detail hub 120 of **Figure 1a**. According to an embodiment, hub 120 of the present invention includes a host-controller-hub (or high-speed hub) controller 180, a hub-agent (or low-speed) hub controller 181, a memory 182, a control unit 183, a clock 183a, a repeater 184, a routing logic 185, and ports 186-189. The controller 180 performs split transactions between hub 120 and host controller 110. Whenever a start-split transaction is received by controller 180 from the host controller 110, controller 180 records the current microframe number in memory 182 as a timestamp for the transaction. An alternate embodiment would record the microframe number as a timestamp in the state records to separate one group of transactions received by the hub during one microframe from those transactions received during another microframe. Controller 181 performs transfers between the hub and one or more agents. The transfers performed by controller 181 are received during start split transactions performed between controller 180 and host controller 110. Memory 182 is coupled to both high-speed hub controller 180 as well as low-speed hub controller 181. Memory includes a pipeline with a

plurality of stages. According to an embodiment, the pipeline has five stages. Each stage of the pipeline corresponds to a microframe as defined above. Each stage of the pipeline has a plurality of transaction states (or transaction records). According to an embodiment, a stage of the pipeline has 19 or fewer transaction states or records. Each stage stores records 5 representative of transfers to be performed. A record may have several fields, such as device and endpoint address, direction of transfer, transfer type, status, and data reference. Status indicates whether the transfer is not ready (prevent the low-speed hub controller from performing it), pending (waiting for execution), pending but must be delayed (because it is an out transaction that spans more than one microframe between the host controller and the hub), 10 or ready (has been performed by the low-speed hub controller). The data reference is a pointer to a starting address in memory for data received (i.e., in transfer) from the agent or data to be sent to the agent (e.g., out transfer).

Hub 120 includes control unit 183 and clock 183a coupled to control unit 183. According to an embodiment, clock 183a generates a microframe indication; i.e., that one-eighth of the time-duration of the classic frame has elapsed. Clock 183a also includes a byte counter 183b that indicates the number of bytes remaining before the beginning of a next 15 microframe. Byte counter 183b is used to control the time when out transfers that span microframes are performed as described above. Control unit 183 is coupled to memory 182 and monitors the records in the stages and prevents the performance of a transfer by the low-speed hub controller 181 if the time at which the transfer is to begin execution at the low-speed hub controller 181 is past the time the transfer was scheduled to be performed or to 20 have completed performance. Specifically, according to an embodiment, a record whose

associated time stamp indicates that it was received by controller 180 from host controller 110 one microframe earlier, is set to a pending status to allow the low-speed hub controller 181 to issue the transaction on the classic bus. A record whose associated timestamp indicates that it was received more than three microframes before the current microframe indication but has

5 not yet been performed by low speed hub controller 181 is marked as old to prepare it for a subsequent corresponding complete-split transaction by the high speed hub controller 180. A record whose associated timestamp indicates that it was received more than four microframes before the current microframe indication but has not yet been performed or is currently being performed by low-speed hub controller 181 is aborted and removed from the pipeline. Clock

10 183a is also coupled to the low-speed hub controller 181. Controller 181 sequences in order through the periodic transaction records that are marked as pending. According to an embodiment, controller 181 sequences through the pending records received in the earliest microframe before proceeding to the records received in the next earliest received microframe and so on. According to an embodiment, controller 181 proceeds to the next earliest received

15 microframe when control unit 181 receives a microframe indication. When control unit 181 receives a microframe indication it changes the status of the transfers in the next earliest received microframe from not ready to pending, and flushes out the transfers that are stale as described above. To perform a transfer controller 181 transfers data to routing logic 185 which gives controller 181 access to one of the ports 186-189. Routing logic 185 gives an agent access to either repeater 184 or controller 181 to allow the transfer of data at a high data rate or a low data rate depending on the agent's configured data rate. Repeater 184 is coupled to controller 180 and routing logic 185. Repeater 184 repeats signals received from (or

destined to) controller 180 and destined to (or received from) a high speed agent coupled to one of ports 186-189.

While sequencing through the pending records received in the earliest microframe, controller 181 operates in accordance with the following rules. First, after performing a transfer between controller 181 and the agent, controller 181 will retrieve the next transaction record and perform the next pending isochronous or interrupt transfer. Second if there is pending delayed transfer, controller 181 will wait until the number of bytes before a microframe tick occurs is less than or equal to the pending delayed transfer's size in bytes before performing the pending delayed transfer. Third, if there is no pending isochronous or interrupt transfers, or pending delayed transfers, controller 181 will perform a pending bulk/control transfer. Fourth, if there is no pending bulk/control transfer, controller 181 waits for the high-speed hub controller 180 to indicate a pending bulk/control or isochronous or interrupt transfer.

When the low-speed hub controller performs a transfer a result is stored in the memory 182. The result may be either a handshake or data received from the agent. Transfers which have been performed have their status changed by the low-speed hub controller 181 to ready and their records are made available to the high-speed hub controller 180. When a complete split transaction is received by the high-speed hub controller 180, the high-speed hub controller 180 examines the record of the most recently performed transfer with the same addressing information and sends a response back to the host controller 110 based on the result of the most recently performed transfer. If the complete split transaction is not

inquiring about a transaction with the same addressing information, the high-speed hub controller 180 will respond back with a NYET.

Thus, a method and apparatus for scheduling transfers between a host controller and a hub has been described. Additionally, a method and apparatus for buffering and performing transfers at a hub has been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be evident to one of ordinary skill in the art that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope the invention as set forth in the claims.

Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.